Quantitative Risk Management
Final Project

John Dodson

October 25, 2021

This assignment is not a regular homework. It is worth half of the module grade for my module. It is due in Canvas at 5:30 PM Central Time on November 1, 2021. You are welcome to discuss the project with our teaching assistant, but she will not be grading it.

This is a group project. Your group can include up to four people. Please include all of the names of the people in your group. Every member of the group will receive the same score. If you do not want to work with a group, you are welcome to submit an individual report, which will then be the sole basis for your score. You may be in at most one group.

Please write up your results into a report and submit a PDF version[1]. You may be able to covert directly to PDF from your word processor, or you may be able to "print to PDF" using a PostScript print driver.

I suggest including the full description of the problem in you report, and add a conclusion in which you discuss your results and and implications for future study or application. Ideally, this report might be something that you make available to a potential employer as a research and writing example.

Please include your appropriately documented code in an appendix.

Stress-Test Risk Measures

For a homogeneous loss in risk factors which are the components of an elliptical random vector, $\boldsymbol{X} \sim \mathrm{E}_d(\boldsymbol{\mu}, \Sigma, \psi(\cdot))$, a coherent risk metric

$$\varrho(L) = \varrho\left(m + \boldsymbol{\lambda}'\boldsymbol{X}\right) = m + r_\varrho(\boldsymbol{\lambda})$$

in terms of the fixed loss m and risk factor allocation $\boldsymbol{\lambda}$, can be expressed as the maximum realized loss over a set of risk factor scenarios S_ϱ that is independent of allocation,

$$\varrho(L) = \sup\left\{m + \boldsymbol{\lambda}'\boldsymbol{x} : \boldsymbol{x} \in S_\varrho\right\}$$

which make up the ellipsoid in $\mathbb{R}^d$

$$S_\varrho = \left\{\boldsymbol{x} : (\boldsymbol{x} - \boldsymbol{\mu})'\,\Sigma^{-1}\,(\boldsymbol{x} - \boldsymbol{\mu}) = \varrho(Y)^2\right\}$$

where $Y \sim S_1(\psi(\cdot))$, is a (univariate) random variable with the same characteristic generator as $\boldsymbol{X}$; that is $\mathrm{E}\, e^{itY} = \psi\left(t^2\right)$.

[1] My edition of The Document Foundation LibreOffice 7.1 can recognize most word processing formats, but sometimes the formatting, equations, or exhibits are altered or corrupted.

Finite Scenario Set Approximation

In practice, we may be interested in approximating the risk factor scenario set by some finite subset. This can be especially useful when the risk factors are highly concordant and when transparency and computational efficiency are important[2].

Market Risk of a Treasury Portfolio

Let's consider the one-day $\alpha = 95\%$ confidence value-at-risk measure, $\varrho(L) = F_L^{\leftarrow}(\alpha)$ where $\boldsymbol{X}$ is a model for the one-day mark-to-market loss on Treasury bonds using the equilibrium calibration from the exercise from October 11, 2021.

The statistical model we fit in that exercise is equivalent to a multivariate Student's-t_4,

$$\boldsymbol{X} \sim \mathrm{E}_8\left(-\hat{\boldsymbol{\mu}}, \hat{\Sigma}, \psi(\cdot)\right)$$

where the characteristic generator is

$$\psi\left(t^2\right) = t^2 K_2\left(\sqrt{2\,t^2}\right)$$

in terms of the modified Bessel function of the second kind. Note that this is the standard, not conventional, version for Student's-t_4; hence the distribution function of Y is

$$F_Y(y) = \frac{1}{2}\left(1 + \frac{y\left(3 + y^2\right)}{\left(2 + y^2\right)^{3/2}}\right)$$

and the diagonals of $\hat{\Sigma}$ are the unconditional variances of the daily returns. The components of $\hat{\boldsymbol{\mu}}$ are identically the secured overnight funding cost per unit notional.

We are interested in approximating value-at-risk with a finite risk factor stress-test set, where the (fixed) scenarios are chosen from the ellipsoid S_ϱ.

$$\varrho_n(L) = \max\left\{m + \boldsymbol{\lambda}'\boldsymbol{x} : \boldsymbol{x} \in S_{\varrho_n}\right\}$$

where

$$\begin{aligned} S_{\varrho_n} &\subset S_\varrho \\ |S_{\varrho_n}| &= n \end{aligned}$$

Problems

1. Calculate $\varrho\left(L_i\right)$, the one-day 95%-confidence equilibrium value-at-risk, for a margined long position in each of the constant-maturity Treasuries $i = 1, 2, \ldots, 8$ (per unit of notional)? **(30 points)**
2. Select $n = 6$ representative stress-test scenarios in S_ϱ to define ϱ_n. **(40 points)**
3. Calculate the eight $\varrho_n\left(L_i\right)$ based on your scenarios. **(10 points)**

[2] The CME SPAN methodology for collateral requirements on futures margin accounts is a variant of this.

Grading Rubric

Twenty out of one hundred points will be based on the follow criteria:

- Your report is clear and professional, and includes an introduction and conclusion. **(5 points)**
- Your narrative and derivations are well-reasoned and convincing. **(5 points)**
- You include or attach your scripts and they are documented. **(5 points)**
- You include appropriate citations and acknowledgements. **(5 points)**

Solution

For affine losses in elliptic risk factors, coherent risk measures for loss random variable L can be expressed in the form

$$\varrho(L) = \mathrm{E}\, L + k_\varrho \sqrt{\operatorname{var} L}$$

Here the risk measure is value-at-risk, so k_ϱ is just a quantile of the (standardized) residual. From the given density, this is a solution to the polynomial

$$y^2 \left(3 + y^2\right)^2 = (2\alpha - 1)^2 \left(2 + y^2\right)^3$$

for the confidence α. This has an exact solution, but deriving it is not instructive. A numerical root finder using a method such as bisection is efficient and accurate in this setting. For $\alpha = 0.95$ I got $k_\varrho \approx 1.507$.

For margined positions, the expected loss is zero, so the value-at-risk is simply k_ϱ times the standard deviation. Using the unconditional variances from the October 11 solution, I got that the results in Table 1.

tenor	j	$\varrho(L_j)$
1Y	1	0.00036
2Y	2	0.0011
3Y	3	0.0016
5Y	4	0.0028
7Y	5	0.0041
10Y	6	0.0057
20Y	7	0.011
30Y	8	0.015

Table 1: 95% value-at-risk per unit notional for margined long positions in par Treasuries

Stress-test scenarios based on principal components

In the scenario version of $r_\varrho(\boldsymbol{\lambda})$, the scenarios $\boldsymbol{x}$ are defined by

$$(\boldsymbol{x} - \boldsymbol{\mu})' \Sigma^{-1} (\boldsymbol{x} - \boldsymbol{\mu}) = \varrho(Y)^2 \tag{1}$$

In the finite approximation, let's restrict to the principal components, which are the eigenvectors corresponding to the largest eigenvalues of Σ. Consider the eigensystem of positive-definite Σ,

$$\Sigma = U' \Lambda U$$

where the columns of U are the eigenvectors, and Λ is a diagonal matrix consisting of the eigenvalues. If we assume that the columns of U are standardized to unit length, then $U'U = I$, which means that $U^{-1} = U'$. Since we need to work with Σ^{-1}, it is useful to note that

$$\begin{aligned}\Sigma^{-1} &= \left(U'\Lambda U\right)^{-1} \\ &= U^{-1}\Lambda^{-1}(U')^{-1} \\ &= U'\Lambda^{-1}U\end{aligned}$$

So the columns of U are also eigenvectors of Σ^{-1}, with the corresponding eigenvalues being just the reciprocals.

Consider a standardized eigenvector $\boldsymbol{v}$ (with $\boldsymbol{v}'\boldsymbol{v} = 1$) and corresponding eigenvalue λ (with $\lambda > 0$), so $\Sigma\boldsymbol{v} = \lambda\boldsymbol{v}$ or $\Sigma^{-1}\boldsymbol{v} = \lambda^{-1}\boldsymbol{v}$. Note that since $\operatorname{E} Y = 0$ and $\operatorname{var} Y = 1$, $\varrho(Y) = k_\varrho$. So, if we let $\boldsymbol{x} - \boldsymbol{\mu} = a\boldsymbol{v}$ for some constant $a \in \mathbb{R}$, (1) translates to

$$(a\boldsymbol{v})'\lambda^{-1}(a\boldsymbol{v}) = k_\varrho^2$$

or $a^2\lambda^{-1} = k_\varrho^2$. Hence, there are two possible stress-test scenarios based on $\boldsymbol{v}$, given by

$$\boldsymbol{x} = \boldsymbol{\mu} \pm k_\varrho\sqrt{\lambda}\boldsymbol{v} \tag{2}$$

Model calibration

In the October 11 assignment, we modeled the dispersion of the standardized daily returns of constant-maturity Treasuries (`disp`). We can interpret this dispersion as an equilibrium covariance if we re-scale the rows and columns so that the diagonal matches the unconditional variance estimates (σ).

```
"correlation"
P̂ = disp ./ sqrt.(diag(disp) .* diag(disp)')

"unconditional covariance"
Σ̂ = P̂ .* (σ * σ')
```

The average daily SOFR funding cost during the hisotrical period was `μ̂ = 0.000047`, just under half a bp/day[3]. We can interpret the expected loss on the components of $\boldsymbol{X}$ as $\boldsymbol{\mu} = -\hat{\mu}\mathbf{1}$. We can also interpret the fixed loss on a margined portfolio of Treasuries as $m = \hat{\mu}\boldsymbol{\lambda}'\mathbf{1}$.

Stress Scenarios

In the October 11 assignment we analyzed the spectrum of the standardized returns—effectively the correlations between Treasury bonds—and we saw support for the three-factor model. In this covariance setting, the largest eigenvector is even more significant, at 96% of the trace; and the largest three amount to 99.7%.

Evaluating (2) for the three largest eigenvectors of $\hat{\Sigma}$ yields six scenarios, which are tabulated in Table 2.

[3] Funding rates prevailing at the end of the historical period were lower than the average.

tenor	first		second		third	
1Y	-0.00019	$+0.00009$	$+0.00007$	-0.00016	$+0.00001$	-0.00010
2Y	-0.00069	$+0.00060$	$+0.00053$	-0.00063	$+0.00019$	-0.00029
3Y	-0.00122	$+0.00112$	$+0.00079$	-0.00089	$+0.00025$	-0.00034
5Y	-0.00243	$+0.00234$	$+0.00130$	-0.00139	$+0.00033$	-0.00043
7Y	-0.00377	$+0.00367$	$+0.00144$	-0.00153	$+0.00029$	-0.00038
10Y	-0.00551	$+0.00542$	$+0.00138$	-0.00148	$+0.00016$	-0.00025
20Y	-0.01067	$+0.01057$	$+0.00037$	-0.00047	-0.00139	$+0.00130$
30Y	-0.01451	$+0.01441$	-0.00160	$+0.00150$	$+0.00068$	-0.00077

Table 2: Treasury return stress scenarios approximating 95% value-at-risk

Approximate value-at-risk

95% value-at-risk for a portfolio of Treasuries is approximated by the fixed loss plus the maximum loss on the portfolio under the nine scenarios in Table 2.

For a portfolio consisting of a unit position in a single CMT category, say category j, the components of the exposure vector $\boldsymbol{\lambda}$ are just $\lambda_i = \delta_{ji}$: The j-th component is a one and the others are zero. Therefore, the maximum loss obtained is simply the largest of the j-th components of the stress scenarios.

With our scenarios above, the largest j-th component always occurs in the second column. This scenario corresponds to an approximately parallel upward shift in yields across the term structure.

tenor	j	max $\left\{\boldsymbol{\delta}_j' \boldsymbol{x}\right\}$
1Y	1	$+0.00009$
2Y	2	$+0.00060$
3Y	3	$+0.00112$
5Y	4	$+0.00234$
7Y	5	$+0.00367$
10Y	6	$+0.00542$
20Y	7	$+0.01057$
30Y	8	$+0.01441$

Table 3: Maximum stress-test loss on long positions in Treasuries

A margined position carries a fixed financing charge, which we determined above was $m = \hat{\mu}$ per unit of notional in each case. This needs to be added back to obtain the approximate risk metric.

The results are presented in Table 4, which repeated the exact results above for comparison.

Discussion

The stress-test approximation based on principal components seems to show good agreement for portfolios consisting of a single long position in a Treasury, with the exception of Treasuries in the 1Y and the 2Y categories, where value-at-risk is substantially underestimated. It would seem that there is a high proportion of idiosyncratic risk in these shortest tenors; and, since their absolute risk is so low, that risk is not captured in the principal components.

tenor	j	$\varrho_6(L_j)$	$\varrho(L_j)$
1Y	1	0.00014	0.00036
2Y	2	0.00064	0.00110
3Y	3	0.0012	0.0016
5Y	4	0.0024	0.0028
7Y	5	0.0037	0.0041
10Y	6	0.0055	0.0057
20Y	7	0.011	0.011
30Y	8	0.014	0.015

Table 4: Approximate & exact 95% daily value-at-risk on long positions in Treasuries

Note that in all cases the stress-test version under-estimates the exact risk. This is a typical result for stress-test risk metrics, and is due to the fact that independent contributions to portfolio variance are being omitted.

If your goal is to estimate value-at-risk, this bias might be considered a weakness. But if the value-at-risk benchmark is not vital in your setting, do not ignore the fact that the stress-test risk metric is a legitimate, coherent measure, which has the advantage of being relatively easy to calculate, communicate, and explain.